

Construction of alternative hypotheses for randomization tests with ordinal outcomes

Jiannan Lu, Peng Ding and Tirthankar Dasgupta*

Harvard University

Abstract

For ordinal outcomes, we construct sequences of alternative hypotheses in increasing departures from the sharp null hypothesis of zero treatment effect on each experimental unit, to help assess the powers of randomization tests in randomized treatment-control experiments.

Keywords: Completely randomized experiment; Potential outcome; Power; Sharp null hypothesis

1. INTRODUCTION

Introduced by Fisher (1935), randomization tests are useful tools for causal inference, because they assess the statistical significance of treatment effects without making any assumptions about the underlying distribution of the outcome. Early theories on randomization tests were developed by Pitman (1938) and Kempthorne (1952), which showed that many statistical procedures can be viewed as approximations of randomization tests. To quote Bradley (1968), “[a] corresponding parametric test is valid only to the extent that it results in the same statistical decision [as the randomization test].” A crucial advantage of randomization tests is their abilities to handle non-standard (e.g., ordinal) outcomes. However, there appears to be limited research on how to assess the powers of randomization tests for ordinal outcomes.

Several researchers like Miller (2006), have pointed out that in many randomized experiments, experimental units cannot be viewed as a random sample drawn from a hypothetical population.

*Address for correspondence: Tirthankar Dasgupta, Department of Statistics, Harvard University, 1 Oxford Street, Cambridge, 02138 Massachusetts, U.S.A. Email: dasgupta@stat.harvard.edu

Therefore it is important to restrict the scope of inference to the finite population of experimental units. The potential outcomes framework (Neyman 1923; Rubin 1974) makes randomization tests easy to interpret, and more importantly, helps us recognize its role in making finite-population inference. However, it does not naturally permit the assessment of the powers of randomization tests, which requires constructing alternatives to the sharp null hypothesis of zero treatment effect on each experimental unit. The existing literature (e.g., Lehmann 1975; Rosenbaum 2010) assesses the powers of randomization tests by invoking infinite population models, primarily to circumvent the difficulty associated with construction of finite population alternatives. Such a construction requires specifying the potential outcomes for all experimental units, and is thus considered “a thankless task” by experts (Rosenbaum 2010). Such a “thankless task” can be made easy by invoking the assumption of independence of potential outcomes, as in some existing literature (e.g., Cheng 2009; Agresti 2010). However, the impact of association between potential outcomes on the powers of randomization tests in a finite population setting has never been investigated before.

In this paper, we demonstrate that it is indeed possible to construct alternative populations of ordinal potential outcomes without invoking the independence assumption. We propose a procedure to construct alternative hypotheses for ordinal outcomes, which is particularly useful in the finite population setting, but is also applicable to infinite populations. Moreover, unlike the existing literature (e.g., Cheng 2009; Agresti 2010) which assume independent potential outcomes, our construction procedure takes into account the dependence structure of the potential outcomes and demonstrates (through simulation studies) that the association indeed does affect the powers of randomization tests in a finite population setting.

The paper proceeds as follows. Section 2 reviews randomization tests of the sharp null hypothesis for ordinal outcomes. Section 3 introduces two measures quantifying departures from the sharp null hypothesis, discusses their relationships to the powers of randomization tests, and proposes a procedure to construct alternative hypotheses in closed forms. Section 4 reports the results of a simulation study that demonstrates how to use the proposed construction procedure to assess the powers of randomization tests. Section 5 presents some concluding remarks.

2. RANDOMIZATION TESTS FOR ORDINAL OUTCOMES

2.1. Potential Outcomes, Sharp Null Hypothesis and Randomization Test

We consider a completely randomized experiment with N units, a binary treatment and an ordinal outcome with J categories labeled as $0, \dots, J-1$, where 0 and $J-1$ are the “worst” and “best” categories, respectively. Under the Stable Unit Treatment Value Assumption (Rubin 1980) that there is only one version of the treatment and no interference among units, we define the pair $\{Y_i(1), Y_i(0)\}$ as the potential outcomes of the i th unit under treatment and control. Let

$$p_{kl} = \text{pr}\{Y_i(1) = k, Y_i(0) = l\} \quad (k, l = 0, 1, \dots, J-1)$$

be the probability of potential outcomes k and l , under treatment and control. Here, the proportion or probability notation $\text{pr}(\cdot)$ can be defined for a finite population with N units, or for an infinite population. The $J \times J$ probability matrix $\mathbf{P} = (p_{kl})_{0 \leq k, l \leq J-1}$ summarizes the joint distribution of the potential outcomes, and plays a crucial role in our later discussion. Let

$$p_{k+} = \sum_{l'=0}^{J-1} p_{kl'}, \quad p_{+l} = \sum_{k'=0}^{J-1} p_{k'l} \quad (k, l = 0, 1, \dots, J-1).$$

The vectors $\mathbf{p}_1 = (p_{0+}, \dots, p_{J-1,+})^T$ and $\mathbf{p}_0 = (p_{+0}, \dots, p_{+,J-1})^T$ characterize the marginal distributions of the potential outcomes under treatment and control.

Using the potential outcomes, we express the sharp null hypothesis as $Y_i(0) = Y_i(1)$ for all i . Under the sharp null hypothesis, the probability matrix \mathbf{P} is diagonal with $p_{j+} = p_{jj} = p_{+j}$, for all $j = 0, 1, \dots, J-1$. To test the sharp null hypothesis, we use data from completely randomized experiments with N_1 units assigned to treatment. For the i th unit, we denote its treatment indicator as W_i , and its observed outcome is consequently $Y_i^{\text{obs}} = W_i Y_i(1) + (1 - W_i) Y_i(0)$. For each j , let n_{0j} and n_{1j} respectively represent the numbers of units exposed to control and treatment with observed outcome j . Given the observed data, we first choose a suitable test statistic, typically a “measure of extremeness” (Brillinger et al. 1978), and then obtain a p -value by comparing the observed value of the test statistic to its randomization distribution.

3. CONSTRUCTION OF ALTERNATIVE HYPOTHESES

To evaluate the powers of randomization tests, we need to construct alternatives to the sharp null hypothesis. We can violate the sharp null hypothesis in two distinct ways:

1. different marginal probabilities, i.e., $\mathbf{p}_1 \neq \mathbf{p}_0$;
2. identical marginal probabilities, but nonzero off diagonal elements in \mathbf{P} .

For example, consider the following probability matrices

$$\mathbf{P}_1 = \begin{pmatrix} \frac{1}{6} & \frac{1}{6} & \frac{1}{6} \\ 0 & \frac{1}{6} & \frac{1}{6} \\ \frac{1}{6} & 0 & 0 \end{pmatrix}, \quad \mathbf{P}_2 = \begin{pmatrix} 0 & \frac{1}{6} & \frac{1}{6} \\ \frac{1}{6} & 0 & \frac{1}{6} \\ \frac{1}{6} & \frac{1}{6} & 0 \end{pmatrix}, \quad \mathbf{P}_3 = \begin{pmatrix} \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \\ \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \\ \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \end{pmatrix}, \quad (1)$$

all of which violate the sharp null hypothesis. In particular, \mathbf{P}_1 has different marginal probabilities; \mathbf{P}_2 and \mathbf{P}_3 have identical marginal probabilities but nonzero off diagonal elements.

Inspired by this example, to construct alternative hypotheses we introduce two measures quantifying violations of the sharp null hypothesis. We use the Hellinger distance τ_{HD} (Hellinger 1909) to quantify the difference between the marginal probabilities:

$$\tau_{HD}(\mathbf{p}_1, \mathbf{p}_0) = \left\{ \frac{1}{2} \sum_{j=0}^{J-1} \left(p_{j+}^{1/2} - p_{+j}^{1/2} \right)^2 \right\}^{1/2}.$$

Other choices include the Kullback–Leibler divergence and total variance distance. Under the sharp null hypothesis $\tau_{HD} = 0$, and therefore nonzero τ_{HD} implies violations of the sharp null hypothesis. However, τ_{HD} relies solely on the marginal probabilities, ignoring the joint distribution of the potential outcomes. For example, the probability matrices \mathbf{P}_2 and \mathbf{P}_3 in (1) violate the sharp null hypothesis although $\tau_{HD} = 0$. To address this issue we use Cohen’s Kappa (Cohen 1960):

$$\kappa(\mathbf{P}) = \left\{ \text{tr}(\mathbf{P}) - \mathbf{p}_1^T \mathbf{p}_0 \right\} / \left(1 - \mathbf{p}_1^T \mathbf{p}_0 \right), \quad (2)$$

where $\text{tr}(\cdot)$ is the trace function. Cohen’s κ relies on the probability matrix \mathbf{P} , and under the sharp null hypothesis $\kappa = 1$ because \mathbf{P} is diagonal.

We now construct sequences of alternative hypotheses by varying the Hellinger distance τ_{HD} and Cohen's κ . To be more specific, we follow a two-step procedure:

1. construct a sequence of marginal probabilities, in an increasing order of τ_{HD} ;
2. for fixed marginal probabilities, construct a sequence of probability matrices in an increasing order of κ , which involves the following sub-steps:

(a) minimize and maximize κ subject to the following constraints:

$$\sum_{k'=0}^{J-1} p_{k'l} = p_{+l}, \quad \sum_{l'=0}^{J-1} p_{kl'} = p_{k+}, \quad p_{kl} \geq 0 \quad (k, l = 0, 1, \dots, J-1);$$

(b) use a convex combination of the minimizer and maximizer to construct probability matrices with intermediate values of κ .

Step 1 helps access the impact of τ_{HD} on the powers of randomization tests, and Step 2 further helps access the impact of κ . For fixed marginal probabilities, Sub-step (a) studies the two extreme cases of “most” and “least” violations of the sharp null hypothesis, and Sub-step (b) addresses the “in between” cases. Therefore, this procedure provides a relatively complete picture of violations of the sharp null hypothesis.

For given marginal probabilities \mathbf{p}_1 and \mathbf{p}_0 , the minimization problem in the above procedure is somewhat intuitive. Consider the probability matrix \mathbf{P}_I with independent potential outcomes, i.e., $p_{kl} = p_{k+}p_{+l}$ for all k, l . If we are not interested in distributions with negatively associated potential outcomes, \mathbf{P}_I minimizes κ as zero. The maximization problem is, however, non-trivial, and almost intractable unless some restrictions are imposed on \mathbf{P} . For the purpose of simplification, we make the following assumption:

Assumption 1. (Stochastic Dominance) For all $j = 1, \dots, J-1$, $\sum_{k=j}^{J-1} p_{k+} \geq \sum_{l=j}^{J-1} p_{+l}$.

As an illustration, consider the three probability matrices in (1). Among these matrices, \mathbf{P}_1 does not satisfy stochastic dominance. Besides the advantage of reducing the number of possible alternative hypotheses, in applied research the stochastic dominance pattern occurs frequently (e.g., Bradley et al. 1962; Bajorski and Petkau 1999), and is termed “positive distributional causal effect” in Ju and Geng (2010). Because of the aforementioned technical convenience and the practical

importance, we first focus on marginal probabilities \mathbf{p}_1 and \mathbf{p}_0 that satisfy Stochastic Dominance, and then discuss general marginal probabilities.

We further simplify the maximization problem by restricting the maximizer to be lower triangular by utilizing a well-known result that for any marginal probabilities satisfying Stochastic Dominance, there exists a probability matrix that is lower triangular. This existence result is a special case of Strassen's theorem (Strassen 1965; Lindvall 1992), which was utilized by Rosenbaum (2001). We are now in a position to state and prove the following theorem that provides the maximum value of κ , and the maximizer itself:

Theorem 1. For any $J \geq 2$, given marginal probabilities \mathbf{p}_1 and \mathbf{p}_0 satisfying Stochastic Dominance, there exists a lower triangular probability matrix \mathbf{P}_+ achieving the upper bound of κ :

$$\kappa(\mathbf{P}) \leq \kappa(\mathbf{P}_+) = \left\{ \sum_{k=0}^{J-1} \min(p_{k+}, p_{+k}) - \mathbf{p}_1^T \mathbf{p}_0 \right\} / (1 - \mathbf{p}_1^T \mathbf{p}_0). \quad (3)$$

Proof of Theorem 1. The proof consists of two parts. We first show that (3) is an upper bound of κ , and then construct a probability matrix attaining it.

For all $k = 0, \dots, J-1$, the diagonal element p_{kk} of matrix \mathbf{P} cannot be greater than either p_{k+} or p_{+k} , i.e., $p_{kk} \leq \min(p_{k+}, p_{+k})$, which implies

$$\text{tr}(\mathbf{P}) \leq \sum_{k=0}^{J-1} \min(p_{k+}, p_{+k}). \quad (4)$$

Substituting (4) in (2) yields the upper bound of κ in (3).

We then sequentially construct a $J \times J$ lower triangular matrix with fixed row sums \mathbf{p}_1 and column sums \mathbf{p}_0 . We start with the last column and proceed backwards. At any point in the construction, we denote the element in the matrix already filled by \tilde{p}_{kl} and those have not by p_{kl} . First, for the last column with index $J-1$, only the last entry needs to be filled, and we set it equal to the corresponding column sum, i.e., $\tilde{p}_{J-1, J-1} = p_{+, J-1}$. Next, for all $r = 1, \dots, J-1$, given all elements in the last r columns are already filled, we consider the problem of filling in the elements of column with index $l = J-r-1$, as shown in Table 1. At this point, the already filled elements in the matrix are \tilde{p}_{kl} , where $k = 0, \dots, J-1$ and $l = J-r, \dots, J-1$.

To fill the column with index $l = J-r-1$, note that all entries for $k < J-r-1$ will be equal

Table 1: Filling in the column with index $l = J - r - 1$ when the last r columns are already filled

Row index (k)	Column index (l)						Row Sum p_{k+}
	0	\dots	$J - r - 1$	$J - r$	\dots	$J - 1$	
0	p_{00}	\dots	0	0	\dots	0	p_{0+}
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
$J - r - 1$	$p_{J-r-1,0}$	\dots	$\min(p_{J-r-1,+}, p_{+,J-r-1})$	0	\dots	0	$p_{J-r-1,+}$
$J - r$	$p_{J-r,0}$	\dots	$p_{J-r,J-r-1} = ?$	$\tilde{p}_{J-r,J-r}$	\dots	0	$p_{J-r,+}$
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
$J - 1$	$p_{J-1,0}$	\dots	$p_{J-1,J-r-1} = ?$	$\tilde{p}_{J-1,J-r}$	\dots	$\tilde{p}_{J-1,J-1} = p_{+,J-1}$	$p_{J-1,+}$
Column sum	p_{+0}	\dots	$p_{+,J-r-1}$	$p_{+,J-r}$	\dots	$p_{+,J-1}$	1

to zero. We set the diagonal element with row index $k = l = J - r - 1$ to be the minimum of the corresponding row and column sums, i.e., $\tilde{p}_{J-r-1,J-r-1} = \min(p_{J-r-1,+}, p_{+,J-r-1})$. Now the difference $p_{+,J-r-1} - \min(p_{J-r-1,+}, p_{+,J-r-1})$ needs to be distributed over the remaining entries of the column. Note that this difference is zero if $\min(p_{J-r-1,+}, p_{+,J-r-1}) = p_{+,J-r-1}$. Therefore, for all $k = J - r, \dots, J - 1$, we make the entry $p_{k,J-r-1}$ proportional to the “remaining balance” $p_{+,J-r-1} - \min(p_{J-r-1,+}, p_{+,J-r-1})$, where we choose the proportionality constant as

$$\frac{\sum_{l=0}^{J-r-1} p_{kl}}{\sum_{k' \geq J-r} \sum_{l=0}^{J-r-1} p_{k'l}}, \quad (5)$$

that is, the ratio of the sum of empty entries in the row with label k and the sum of empty entries in all rows below the one labeled $J - r - 1$. Both the numerator and denominator of (5) can be expressed in terms of the given marginals and the already filled-in entries in the last r columns:

$$\sum_{l=0}^{J-r-1} p_{kl} = p_{k+} - \sum_{l \geq J-r} \tilde{p}_{kl}, \quad \sum_{k' \geq J-r} \sum_{l=0}^{J-r-1} p_{k'l} = \sum_{k' \geq J-r} \left(p_{k'+} - \sum_{l \geq J-r} \tilde{p}_{k'l} \right) \quad (6)$$

and hence can be computed uniquely. The construction method, (5) or (6), eventually leads to the following iterative imputation equation for all $k = J - r, \dots, J - 1$:

$$\tilde{p}_{k,J-r-1} = \frac{p_{k+} - \sum_{l \geq J-r} \tilde{p}_{kl}}{\sum_{k' \geq J-r} \left(p_{k'+} - \sum_{l \geq J-r} \tilde{p}_{k'l} \right)} \{p_{+,J-r-1} - \min(p_{J-r-1,+}, p_{+,J-r-1})\}. \quad (7)$$

We need to show the constructed matrix

$$\mathbf{P}_+ = \begin{pmatrix} \tilde{p}_{00} & 0 & \dots & 0 \\ \tilde{p}_{10} & \tilde{p}_{11} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{p}_{J-1,0} & \tilde{p}_{J-1,1} & \dots & \tilde{p}_{J-1,J-1} \end{pmatrix}$$

indeed satisfies (i) $\tilde{p}_{kl} \geq 0$ for all $k, l = 0, \dots, J-1$, (ii) the equality in condition (4), for which a sufficient condition is $\tilde{p}_{kk} = \min(p_{k+}, p_{+k})$ for all $k = 0, \dots, J-1$, (iii) the vector of column sums is \mathbf{p}_0 , i.e., $\sum_{k=0}^{J-1} \tilde{p}_{kl} = p_{+l}$ for all $l = 0, \dots, J-1$, and (iv) the vector of row sums is \mathbf{p}_1 , i.e., $\sum_{l=0}^{J-1} \tilde{p}_{kl} = p_{k+}$ for all $k = 0, \dots, J-1$.

Among (i)–(iv) described above, (i)–(iii) follow directly by the construction of \mathbf{P}_+ described above. We need only to prove that $\sum_{l=0}^{J-1} \tilde{p}_{kl} = p_{k+}$ for all $k = 0, \dots, J-1$. By Stochastic Dominance, we have $p_{0+} \leq p_{+0}$, implying that $\tilde{p}_{00} = p_{0+}$. Therefore the sum of the first row of \mathbf{P}_+ is p_{0+} . Now for all $k = 1, \dots, J-1$, by substituting $r = J-1$ in (7), or by filling up the first column given the last $J-1$, we have

$$\begin{aligned} \tilde{p}_{k0} &= \frac{p_{k+} - \sum_{l=1}^{J-1} \tilde{p}_{kl}}{\sum_{k=1}^{J-1} \left(p_{k+} - \sum_{l=1}^{J-1} \tilde{p}_{kl} \right)} (p_{+0} - p_{0+}) = \frac{p_{k+} - \sum_{l=1}^{J-1} \tilde{p}_{kl}}{(1 - p_{0+}) - (1 - \tilde{p}_{00})} (p_{+0} - p_{0+}) \\ &= \frac{p_{k+} - \sum_{l=1}^{J-1} \tilde{p}_{kl}}{p_{+0} - p_{0+}} (p_{+0} - p_{0+}) = p_{k+} - \sum_{l=1}^{J-1} \tilde{p}_{kl}, \end{aligned}$$

which implies that $\sum_{l=0}^{J-1} \tilde{p}_{kl} = p_{k+}$. The proof is complete. \square

In the above proof of Theorem 1, we suggest a way to constructing the maximizer \mathbf{P}_+ . Next we discuss the uniqueness of \mathbf{P}_+ . By restricting \mathbf{P}_+ to be lower triangular and its $(j+1)$ th diagonal element p_{jj} to be $\min(p_{j+}, p_{+j})$, what remain to be determined are the $(J-1)J/2$ off diagonal elements. Note that there are $(2J-3)$ constraints associated with them. The equality $(J-1)J/2 = 2J-3$ holds if and only if $J = 2$ or 3 .

The case with $J = 2$ corresponds to binary outcomes, which occur frequently in both methodology and applied research. For a recent discussion of finite population inference for binary data, see Ding and Dasgupta (2015). The following corollary provides the maximizer under $J = 2$. Although

it is a special case of Theorem 1, we provide a direct proof to rigorously show the uniqueness of the maximizer.

Corollary 1. For $J = 2$, given marginal probabilities \mathbf{p}_1 and \mathbf{p}_0 that satisfy Stochastic Dominance, the following matrix is the unique maximizer of κ :

$$\mathbf{P}_+ = \begin{pmatrix} p_{0+} & 0 \\ p_{1+} - p_{+1} & p_{+1} \end{pmatrix}. \quad (8)$$

Proof of Corollary 1. Because \mathbf{p}_1 and \mathbf{p}_0 satisfy Stochastic Dominance, we have $p_{0+} \leq p_{+0}$ and $p_{1+} \geq p_{+1}$, implying that the diagonal elements of the maximizer are $p_{00} = p_{0+}$ and $p_{11} = p_{+1}$. Because the row sums of the maximizer are \mathbf{p}_1 , we uniquely determine the entries of the maximizer, as shown in (8). The maximizer has nonnegative entries because $p_{1+} \geq p_{+1}$, and its column sums are \mathbf{p}_0 because $p_{0+} + p_{1+} - p_{+1} = p_{+0}$. The proof is complete. \square

The case with $J = 3$ corresponds to three-level outcomes, which are also important in practice. For example, in a clinical trial we can describe the status of a patient as “deterioration,” “no change” or “improvement” (Bajorski and Petkau 1999). The following corollary gives the maximizer for $J = 3$. Again, we provide a direct proof.

Corollary 2. For $J = 3$, given marginal probabilities \mathbf{p}_1 and \mathbf{p}_0 that satisfy Stochastic Dominance, the following matrix is the unique maximizer of κ :

$$\mathbf{P}_+ = \begin{pmatrix} p_{0+} & 0 & 0 \\ p_{1+} - \min(p_{+1}, p_{1+}) & \min(p_{+1}, p_{1+}) & 0 \\ p_{2+} - p_{+2} - \{p_{+1} - \min(p_{+1}, p_{1+})\} & p_{+1} - \min(p_{+1}, p_{1+}) & p_{+2} \end{pmatrix}. \quad (9)$$

Proof of Corollary 2. Because \mathbf{p}_1 and \mathbf{p}_0 satisfy Stochastic Dominance, we have $p_{0+} \leq p_{+0}$ and $p_{2+} \geq p_{+2}$, which implies that the diagonal elements of the maximizer are $p_{00} = p_{0+}$, $p_{11} = \min(p_{+1}, p_{1+})$, and $p_{22} = p_{+2}$. First, because the first row sum and third column sum are respec-

tively p_{0+} and p_{+2} , the maximizer is in the following form:

$$\mathbf{P}_+ = \begin{pmatrix} p_{0+} & 0 & 0 \\ ? & \min(p_{+1}, p_{1+}) & 0 \\ ? & ? & p_{+2} \end{pmatrix},$$

where “?” denotes an entry yet to be determined. Second, because the second row sum and column sum are respectively p_{1+} and p_{+1} , the maximizer is in the following form:

$$\mathbf{P}_+ = \begin{pmatrix} p_{0+} & 0 & 0 \\ p_{1+} - \min(p_{+1}, p_{1+}) & \min(p_{+1}, p_{1+}) & 0 \\ ? & p_{+1} - \min(p_{+1}, p_{1+}) & p_{+2} \end{pmatrix}.$$

Third, because the third row sum is p_{2+} , we uniquely determine the maximizer, as in (9). Fourth, \mathbf{P}_+ has nonnegative entries, because

$$p_{2+} - p_{+2} - \{p_{+1} - \min(p_{+1}, p_{1+})\} = \min(p_{2+} - p_{+2}, p_{+0} - p_{0+}) \geq 0.$$

Finally, \mathbf{P}_+ has row sums \mathbf{p}_1 , and column sums \mathbf{p}_0 , because

$$p_{0+} + p_{1+} - \min(p_{+1}, p_{1+}) + p_{2+} - p_{+2} - \{p_{+1} - \min(p_{+1}, p_{1+})\} = p_{+0}.$$

The proof is complete. □

We end this section by constructing probability matrices with intermediate values of κ . Given the minimizer \mathbf{P}_I and maximizer \mathbf{P}_+ , let $\mathbf{P}_\lambda = \lambda \mathbf{P}_I + (1 - \lambda) \mathbf{P}_+$. We view $\lambda \in [0, 1]$ as a sensitivity parameter, because we cannot estimate it from the observed data. The resulting probability matrices have the same marginal probabilities as \mathbf{P}_I and \mathbf{P}_+ , and subsequently the same Hellinger distances. However, they have different κ depending on λ because $\kappa(\mathbf{P}_\lambda) = (1 - \lambda)\kappa(\mathbf{P}_+)$. For the infinite population framework, our constructed sequence of alternative hypotheses are thus $\{\mathbf{P}_\lambda\}_{\lambda \in [0, 1]}$. For the finite population framework, because any entry of a well-defined probability

matrix \mathbf{P} is multiples of $1/N$, we propose a calibration step by letting

$$\tilde{p}_{kl}(\lambda) = \begin{cases} \frac{\lfloor Np_{kl}(\lambda) \rfloor}{N} & \text{if } k \neq l, \\ p_{+l} - \sum_{k' \neq l} \frac{\lfloor Np_{k'l}(\lambda) \rfloor}{N} & \text{if } k = l, \end{cases}$$

where $\lfloor \cdot \rfloor$ is the floor function. By definition, the column sums of $\tilde{\mathbf{P}}_\lambda$ are \mathbf{p}_0 . Let $\Lambda(\mathbf{p}_1, \mathbf{p}_0)$ denote the set containing all λ 's such that the row sums of $\tilde{\mathbf{P}}_\lambda$ are \mathbf{p}_1 , and our constructed sequence of alternative hypotheses are therefore $\{\tilde{\mathbf{P}}_\lambda\}_{\lambda \in \Lambda(\mathbf{p}_1, \mathbf{p}_0)}$. In practice, we can use a grid search to obtain an approximation of $\Lambda(\mathbf{p}_1, \mathbf{p}_0)$.

4. A SIMULATION STUDY

We demonstrate how the above construction facilitates assessment of powers of randomization tests. We use the squared Mann–Whitney U -statistic (Agresti 2002):

$$U^2 = \left[\sum_{k=0}^{J-1} \sum_{l=0}^{J-1} n_{1k} n_{0l} \{I(k > l) - I(k < l)\} \right]^2.$$

Another commonly-used test statistic for categorical data is the χ^2 -statistic. However, we do not use it in this paper, because it does not utilize the order information and therefore is less powerful.

Although closed-form expressions of the powers of randomization test using the U^2 statistic are difficult to obtain, numerical calculations by Monte Carlo are straightforward, once we determine the alternative hypothesis \mathbf{P} . We will follow three steps:

1. under \mathbf{P} generate 2×10^5 independent treatment assignments and obtain the corresponding observed data sets;
2. for each observed data set calculate the p -value of the randomization test using the observed value of U^2 and its simulated null distribution;
3. approximate the power of the U^2 statistic as the proportion of the p -values that are smaller than the significance level $\alpha = 0.05$.

In this simulation study, we construct alternative hypotheses using the following four sets of marginal probabilities, two with $J = 2$ and two with $J = 3$:

1. $\mathbf{p}_1 = (3/10, 7/10)^T$, $\mathbf{p}_0 = (3/5, 2/5)^T$, $\tau_{HD} = 0.216$;
2. $\mathbf{p}_1 = (1/2, 1/2)^T$, $\mathbf{p}_0 = (4/5, 1/5)^T$, $\tau_{HD} = 0.227$;
3. $\mathbf{p}_1 = (1/4, 1/4, 1/2)^T$, $\mathbf{p}_0 = (2/5, 2/5, 1/5)^T$, $\tau_{HD} = 0.227$;
4. $\mathbf{p}_1 = (9/40, 9/40, 11/20)^T$, $\mathbf{p}_0 = (2/5, 2/5, 1/5)^T$, $\tau_{HD} = 0.261$.

For each case, we let the sample sizes $N = 120, 160, 240$, and the sensitivity parameters $\lambda = 0, 1/4, 1/2, 3/4, 1$. We then construct the probability matrices, which share the same marginals. For each probability matrix $\tilde{\mathbf{P}}_\lambda$, we use the aforementioned Monte Carlo procedure to calculate the powers. On one hand, different cases allow us to study the impact of τ_{HD} on the powers. On the other hand, within each case we can study the impact of κ on the powers. The simulation results are summarized in Figure 1, from which we draw the following conclusions. For all fixed sample sizes, the power functions of Case 2 dominate those of Case 1, and the power functions of Case 4 dominate those of Case 3. Therefore, for fixed J the power increases as the Hellinger distance increases. Furthermore, for fixed marginals and sample size, the power increases as κ decreases, or equivalently as λ increases. However, this dependence becomes weaker as the sample size increases, because the power converges to one.

We can use the demonstrated methodology to compare the power functions of different test statistics, and also to determine sample sizes that guarantee a pre-specified power. For instance, in Case 3, we cannot guarantee a power of 0.95 with a sample of size 120, but we can with a sample of size 160.

In summary, for a finite population, the power of the randomization test using U^2 depends on the marginal difference of the potential outcomes as well as the association between them. In particular, the power increases as the marginal difference increases, and given the marginals fixed, the power increases as the association decreases. Furthermore, the power converges to one as the sample size increases, for any case with nonzero marginal difference. The above conclusions appear to confirm our intuition, because it should be easier to reject the sharp null hypothesis given larger differences between the marginals, and given the marginals fixed, it should be easier to reject the sharp null hypothesis given smaller associations between the potential outcomes. Our findings conform to Plackett (1977)'s and Chernoff (2004)'s results about the classical 2×2 tables: the

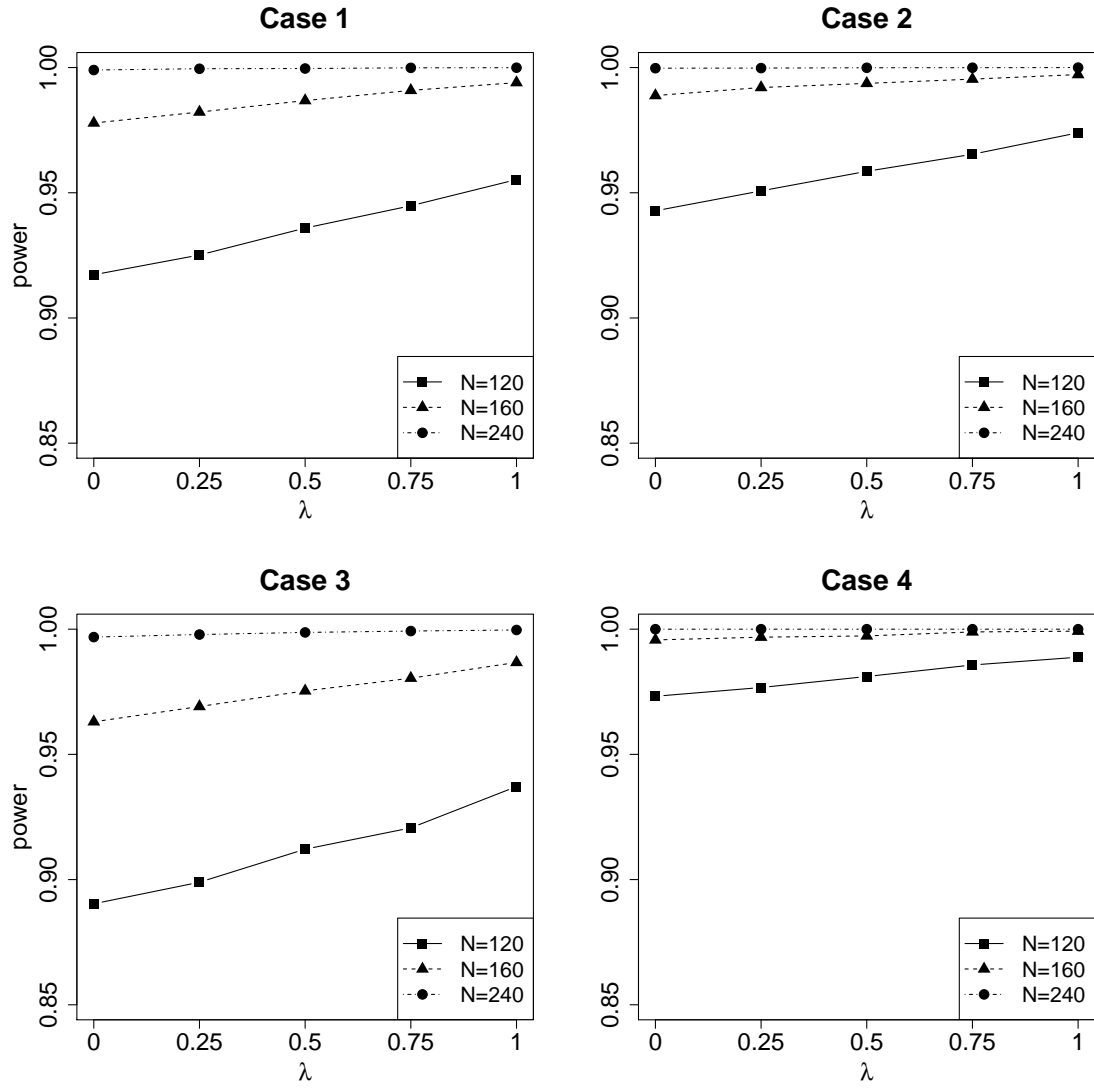


Figure 1: Statistical Powers of Randomization Tests Using U^2 .

marginals of the contingency tables contain limited amount of information about the association with finite samples, which becomes negligible asymptotically.

5. DISCUSSION

In this paper, we construct sequences of finite populations of ordinal outcomes in increasing departures from the sharp null hypothesis of no treatment effect. Our construction procedure is useful for evaluating of the powers of randomization tests. In particular, our construction procedure takes into account the dependence structure of the potential outcomes, whereas existing literature often assume independent potential outcomes. Through a simulation study, we demonstrate that the association between potential outcomes indeed affects the powers of randomization tests. We argue that taking into account the association between potential outcomes is crucial for conducting randomization tests in practice, for example when determining sample sizes that guarantee a pre-specified power.

There are multiple future directions based on our work. First, although we adopt a numerical approach, it is possible to derive the asymptotic distribution of the U^2 statistic under the sharp null hypothesis. Second, we can derive the maximizer of κ for general marginal probabilities that do not satisfy Stochastic Dominance. Third, we can incorporate covariate information to further improve the powers of randomization tests. Fourth, while the Fisherian randomization-based inference is a useful first step, Neymanian and Bayesian counterparts of causal inference for ordinal outcomes are still needed. For some recent developments, see Lu et al. (2015).

ACKNOWLEDGEMENTS

The authors thank a reviewer and the editor for their thoughtful comments that helped clarify several points better and substantially improved the presentation of the paper.

REFERENCES

- Agresti, A. (2002). *Categorical Data Analysis, 2nd Edition*. New York: John Wiley and Sons.
- Agresti, A. (2010). *Analysis of Ordinal Categorical Data, 2nd ed.* Hoboken, New Jersey: John Wiley and Sons.

- Bajorski, P. and Petkau, J. (1999). Nonparametric two-sample comparison of changes on ordinal responses. *Journal of American Statistical Association*, 94:970–978.
- Bradley, J. V. (1968). *Distribution-Free Statistical Tests*. Upper Saddle River, NJ: Prentice Hall.
- Bradley, R. A., Katti, S. K., and Coons, I. J. (1962). Optimal scaling for ordered categories. *Psychometrika*, 27:355–374.
- Brillinger, D. R., Jones, L. V., and Tukey, J. W. (1978). Report of the statistical task force to the weather modification advisory board. *The Management of Weather Resources, II: The Role of Statistics on Weather Resources Management*, Government Printing Office, Washington D.C.
- Cheng, J. (2009). Estimation and inference for the causal effect of receiving treatment on a multinomial outcome. *Biometrics*, 65:96–103.
- Chernoff, H. (2004). Information, for testing the equality of two probabilities, from the margins of the 2×2 table. *Journal of Statistical Planning and Inference*, 121:209–214.
- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20:37–46.
- Ding, P. and Dasgupta, T. (2015). A potential tale of two by two tables from completely randomized experiments. *Journal of American Statistical Association*, in press.
- Fisher, R. A. (1935). *The Design of Experiments, 1st Edition*. Edinburgh, London: Oliver and Boyd.
- Hellinger, E. (1909). Neue begründung der theorie quadratischer formen von unendlichvielen veränderlichen. *Journal für die reine und angewandte Mathematik*, 136:210–271.
- Ju, C. and Geng, Z. (2010). Criteria for surrogate end points based on causal distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72:129–142.
- Kempthorne, O. (1952). *The Design and Analysis of Experiments*. New York: John Wiley and Sons.
- Lehmann, E. L. (1975). *Nonparametrics: Statistical Methods Based on Ranks, 1st Edition*. San Francisco: Holden-Day.

- Lindvall, T. (1992). *Lectures on the Coupling Method*. New York: Wiley.
- Lu, J., Ding, P., and Dasgupta, T. (2015). Sharp bounds of causal effects on ordinal outcomes. *arXiv: 1507.01542*.
- Miller, S. (2006). *Experimental Design and Statistics*. London: Routledge.
- Neyman, J. (1923). On the application of probability theory to agricultural experiments. Essay on principals. Section 9. [Translated in *Statistical Science* (1990) **5**, 465–480].
- Pitman, E. J. (1938). Significance tests which may be applied to samples from any populations: III. The analysis of variance test. *Biometrika*, 29:322–335.
- Plackett, R. L. (1977). The marginal totals of a 2×2 table. *Biometrika*, 64:37–42.
- Rosenbaum, P. R. (2001). Effects attributable to treatment: inference in experiments and observational studies within a discrete pivot. *Biometrika*, 88:219–231.
- Rosenbaum, P. R. (2010). *Design of Observational Studies*. New York: Springer.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66:688–701.
- Rubin, D. B. (1980). Comment on “Randomization analysis of experimental data: the Fisher randomization test” by D. Basu. *Journal of American Statistical Association*, 75:591–593.
- Strassen, V. (1965). The existence of probability measures with given marginals. *Annals of Mathematical Statistics*, 36:423–439.